arXiv:1512.01483v4 [math.CO] 25 Jul 2016

SWEEPING UP ZETA

HUGH THOMAS AND NATHAN WILLIAMS

Abstract. We repurpose the main theorem of [TW14] to prove that modular sweep maps are bijective. We construct the inverse of the modular sweep map by passing through an intermediary set of equitable partitions; motivated by an analogy to stable marriages, we prove that the set of equitable partitions for a fixed word forms a distributive lattice when ordered componentwise. We conclude that the general sweep maps defined in [ALW14b] are bijective. As a special case of particular interest, this gives the first proof that the zeta map on rational Dyck paths is a bijection.

1. Introduction

The sweep map of [ALW14b] is a broad generalization of the zeta map on Dyck paths, originally defined by J. Haglund and M. Haiman in the context of the study of diagonal harmonics. Proving bijectivity of the sweep map was an open problem with significant implications in the study of rational Catalan combinatorics (see Section 2). We solve this problem in Theorem 6.1.

To put the reader in the right frame of mind for our solution, we begin with a "real world" interpretation of the problem.

1.1. **Scheduling Tasks.** Consider the following "real world" problem of scheduling daily tasks (for example, on a computer). The day is divided into m hours, and there are N tasks to be carried out each day (numbered, or prioritized, from 1 to N). For simplicity, let us suppose that the total amount of time required to carry out all of the tasks is some multiple of m. Each task takes an integer number of hours to complete: a task can take zero hours, but no task takes as much as an entire day. Tasks start on the hour, and cannot be interrupted once started. Tasks can be worked on concurrently, and the starting order for the tasks within the day is specified—the jth task must start before or at the same time as the $(j+1)$st task. The schedule repeats every day, so tasks can start at the end of one day and finish at the beginning of the next day.

Under these assumptions, it is reasonable to ask for an assignment of starting hours for the tasks so that the workload throughout the day is constant. We call an answer—the starting times of the tasks—an *equitable partition* for our scheduling problem. It is not immediately clear that an answer necessarily exists.

Example 1.1. Suppose we divide the day into $m = 5$ hours, and we have a list of $N = 7$ tasks such that the length (in hours) that it takes to complete the jth task is the jth element of the sequence $[1, 3, 3, 1, 4, 2, 1]$. Since it would take $15 = 3 \cdot 5$ hours to complete all the tasks, we should work on three tasks at a time.

Starting tasks in the hours $[1, 1, 2, 2, 3, 4, 5]$ (so that tasks one and two both start in hour 1, tasks three and four in hour 2, and so on until task seven starts in hour 5) or in the hours $[1, 2, 2, 4, 5, 5, 5]$ both yield equitable partitions. These are

Date: July 26, 2016.

2000 *Mathematics Subject Classification.* 05E99.

Key words and phrases. zeta map, sweep map, rational Catalan combinatorics, affine Weyl groups, stable marriages.

illustrated below—each task is assigned a row and each hour a column, and we mark which tasks we worked on when with a symbol ■ in the corresponding cell. The conditions for an equitable partition constrain the starting point of the jth row to be weakly to the left of that of the $(j+1)$st row, and force each column to have the same number of copies of ■.

		Hours				
		1	2	3	4	5
Tasks	1	■	·	·	·	·
	2	■	■	■	·	·
	3	·	■	■	■	·
	4	·	■	·	·	·
	5	■	·	■	■	■
	6	·	·	·	■	■
	7	·	·	·	·	■

and

		Hours				
		1	2	3	4	5
Tasks	1	■	·	·	·	·
	2	·	■	■	■	·
	3	·	■	■	■	·
	4	·	·	·	■	·
	5	■	■	■	·	■
	6	■	·	·	·	■
	7	·	·	·	·	■

Now imagine that an (admittedly, somewhat fussy) inspector arrives at the last hour of a day, and wants to watch each of the tasks being done—one at a time, each one performed from beginning to end without interruption. After an observed task completes, the inspector next watches the lowest-numbered task beginning as promptly as possible among those tasks they have not yet watched. A *successful partition* is an equitable partition that allows our inspector to watch each task performed exactly once, without any delay between the end of one task and the beginning of another.

Example 1.2. The equitable partition $[1, 2, 2, 4, 5, 5, 5]$ (above, on the right) is successful because our inspector can observe the tasks without interruption as follows:

		Hours														
		5	1	2	3	4	5	1	2	3	4	5	1	2	3	4
Tasks	1	·	·	·	·	·	·	·	·	·	·	·	■	·	·	·
	2	·	·	·	·	·	·	·	■	■	■	·	·	·	·	·
	3	·	·	·	·	·	·	·	·	·	·	·	·	■	■	■
	4	·	·	·	·	■	·	·	·	·	·	·	·	·	·	·
	5	■	■	■	■	·	·	·	·	·	·	·	·	·	·	·
	6	·	·	·	·	·	■	■	·	·	·	·	·	·	·	·
	7	·	·	·	·	·	·	·	·	·	·	■	·	·	·	·

The equitable partition $[1, 1, 2, 2, 3, 4, 5]$ (above, left) is *not* successful because the inspector experiences delays between the end of the third task and the beginning of the second, and between the end of the sixth task and the beginning of the fourth:

		Hours																
		5	1	2	3	4	5	1	2	3	4	5	1	2	3	4	5	1
Tasks	1	·	■	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
	2	·	·	·	·	·	·	■	■	■	·	·	·	·	·	·	·	·
	3	·	·	■	■	■	·	·	·	·	·	·	·	·	·	·	·	·
	4	·	·	·	·	·	·	·	·	·	·	·	·	■	·	·	·	·
	5	·	·	·	·	·	·	·	·	·	·	·	·	·	■	■	■	■
	6	·	·	·	·	·	·	·	·	·	■	■	·	·	·	·	·	·
	7	■	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·

It turns out that the successful partition is unique—it is the equitable partition for which each task is started as late as possible (among all equitable partitions). In an amusing twist to what "real-life" experience might suggest, this particular inspector can only be satisfied by procrastination.

1.2. **The Modular Sweep Map.** A solution to the scheduling problem above is closely related to inverting a simple and curious map defined on words.

As above, we let $m, N \in \mathbb{N}$, but we now write $\mathcal{A}$ for the set of words of length N on the alphabet $\{0, 1, 2, \ldots, m-1\}$. For a word $\mathsf{w} = \mathsf{w}_1\mathsf{w}_1 \cdots \mathsf{w}_N \in \mathcal{A}$ and for $1 \leq j \leq N$, define the *modular level* of the letter w_j to be $\ell_j := \sum_{i=1}^{j} \mathsf{w}_i \bmod m$.

The *modular sweep map* is the function $\mathsf{sweep}_m : \mathcal{A} \to \mathcal{A}$ that sorts $\mathsf{w} \in \mathcal{A}$ according to its modular levels as follows: initialize $\mathsf{u} = \emptyset$ to be the empty word. For $k = m-1, \ldots, 2, 1, 0$, read w from right to left and append to u all letters w_j whose level ℓ_j is equal to k. Define $\mathsf{sweep}_m(\mathsf{w}) := \mathsf{u}$.

Example 1.3. Let $m = 5$ and $N = 7$. We compute the modular levels of the word $\mathsf{w} = 3113214 \in \mathcal{A}$ by summing the initial letters of w modulo m and obtain the image $\mathsf{u} := \mathsf{sweep}_m(\mathsf{w})$ by sorting according to the levels.

ℓ	3	4	0	3	0	1	0		ℓ	4	3	3	1	0	0	0
w	3	1	1	3	2	1	4	$\xmapsto[\mathsf{sweep}_m]{}$	u	1	3	3	1	4	2	1

Note that u records the lengths of the tasks, in order of priority, while w records the sequence of task lengths in the reverse of the order in which they are watched by the inspector in Example 1.2. Indeed, as we shall make precise, finding a sequence of tasks that the inspector can watch in order with no breaks amounts to inverting the modular sweep map.

Our main result—proven in Section 4.5—is that sweep_m is invertible.[1]

Theorem 1.4. *The modular sweep map is a bijection $\mathcal{A} \to \mathcal{A}$.*

1.3. **Organization.** The remainder of this paper is organized as follows.

We give a brief history in Section 2 by recalling the different contexts in which the modular sweep map has appeared.

In Section 3.2, we define the modular presweep map. This map differs from the modular sweep map in that it preserves the additional information of the modular levels. It is easy to invert the modular presweep map, as described in Section 3.3; partitions for which the inverse modular presweep map concludes are called successful partitions.

In Section 4.1, we introduce the notion of *equitable partitions* and show that a succcessful partition is equitable. Using an algorithm communicated to us by F. Aigner, C. Ceballos, and R. Sulzgruber (inspired by our related Algorithm 5), we construct the rightmost equitable partition in Theorem 4.4, and we show how to modify any equitable partition to produce a successful partition in Theorem 4.10. Theorem 4.12 concludes that the rightmost equitable partition and successful partition are the same.

We apply the results of Sections 3 and 4 to prove Theorem 1.4—that the modular sweep map is a bijection—in Section 4.5.

In analogy to the stable marriage problem of D. Gale and L. Shapley, we construct the leftmost equitable partition in Theorem 5.2, and we show that the set of all equitable partitions may be given the structure of a distributive lattice in

[1] As G. Warrington pointed out to us at the American Institute of Mathematics in 2012—sorting is not usually an invertible operation!

Theorem 5.8. Section 5 is independent of the applications of Theorem 1.4 given in Section 6, and may be skipped.

In Section 6, we use Theorem 1.4 to solve two problems from the literature: we show in Section 6.1 how to invert the sweep map of [ALW14b] on words with letters in $\mathbb{Z}$ (rather than $\mathbb{Z}/m\mathbb{Z}$), and we conclude in Section 6.2 that the zeta map is bijective on Dyck paths and rational Dyck paths.

Acknowledgements

The second author is indebted to Vic Reiner and Dennis Stanton for originally suggesting this area.

We thank Drew Armstrong, Cesar Ceballos, Adriano Garsia, Nick Loehr, and Greg Warrington for their encouragement and enthusiasm, close reading, detection of typos, and suggestions for improvements to the exposition. We thank Marko Thiel for pointing out that the argument of [ALW14b, Proposition 3.2] could be extended to arbitrary alphabets in Theorem 6.3. We warmly thank Florian Aigner, Cesar Ceballos, and Robin Sulzgruber for finding Algorithm 3 in the course of their study of an early draft of this paper, and then patiently explaining it to us at FPSAC 2016.

Both authors thank their parents for introducing them to stable marriages.

This research was supported by NSERC and the Canada Research Chairs program.

2. History

2.1. **Diagonal Harmonics and the Zeta Map.** In their study of the space $\mathcal{DH}_n$ of diagonal harmonics [GH96], A. Garsia and M. Haiman defined a rational function $C_n(q,t)$, symmetric in q and t, with the property that $C_n(1,1) = \frac{1}{n+1}\binom{2n}{n}$. They conjectured that $C_n(q,t)$ was actually a *polynomial* in q and t with nonnegative coefficients—specializing one of the statistics to 1, they gave a combinatorial interpretation of this polynomial using the area statistic on *n-Dyck paths* (lattice paths from $(0,0)$ to (n,n) that stay above the diagonal $y=x$):

$$C_n(q,1) = C_n(1,q) = \sum_{w \text{ an } n\text{-Dyck path}} q^{\text{area}(w)}.$$

The search was on to find a statistic that manifested nonnegativity—an unknown statistic with the property that

$$C_n(q,t) = \sum_{w \text{ an } n\text{-Dyck path}} q^{\text{area}(w)} t^{\text{unknown}(w)}.$$

In [Hag03], "after a prolonged investigation of tables of $C_n(q,t)$," J. Haglund invented the idea of a *bounce path*, which he used to propose exactly such a statistic. A. Garsia and J. Haglund subsequently used these ideas to prove nonnegativity of $C_n(q,t)$ in [GH02].

As the legend goes, A. Garsia sent a cryptic email to M. Haiman announcing J. Haglund's discovery—without providing any specifics as to what the statistic was. Shortly after, M. Haiman announced that he, too, had produced the desired statistic.[2] Remarkably, J. Haglund's statistic and M. Haiman's statistic were *different*. In modern language, J. Haglund's statistic is known as bounce, while M. Haiman's is

[2] A. Garsia subsequently expressed regret that he didn't send this email to M. Haiman several years earlier.

dinv. M. Haiman and J. Haglund quickly developed a bijection from n-Dyck paths to themselves—the *zeta map* ζ (see Section 6.3) [AKOP02, Hag03]—such that

$$(\text{area}(w), \text{bounce}(w)) = (\text{dinv}(\zeta(w)), \text{area}(\zeta(w))).$$

As Dyck paths have been generalized (say, as in Section 6.2, to lattice paths from $(0,0)$ to (a,b) that stay above the main diagonal), so too have these zeta maps [Loe03, Egg03, GM14, LLL14]. A modern perspective is that there is only one statistic—area—along with a generalized zeta map [ALW14b]. If such a zeta map is bijective on a set of generalized Dyck paths $\mathcal{D}$, one can combinatorially define polynomials

$$\mathcal{D}(q,t) := \sum_{w \in \mathcal{D}} q^{\text{area(w)}} t^{\text{area}(\zeta(\text{w}))},$$

so that (by construction) $\mathcal{D}(q,1) = \mathcal{D}(1,q)$. Surprisingly, these polynomials also often happen to be symmetric in q and t.

Proving invertibility of these generalized zeta maps has been a traditionally difficult problem; combinatorially proving (q,t)-symmetry has been intractable. Most recently, D. Armstrong, N. Loehr, and G. Warrington have found a very general versions of the zeta map, which they called *sweep maps* [ALW14b, Section 3.4]. C. Reutenauer independently developed a related map in an unpublished letter to A. Garsia, A. Hicks, and E. Leven [Reu].

2.2. **Suter's Cyclic Symmetry.** In [Sut02], R. Suter defined a striking cyclic symmetry of order $n+1$ on the subposet of Young's lattice consisting of those partitions with largest hook at most n. Figure 1 illustrates this for $n = 4$.

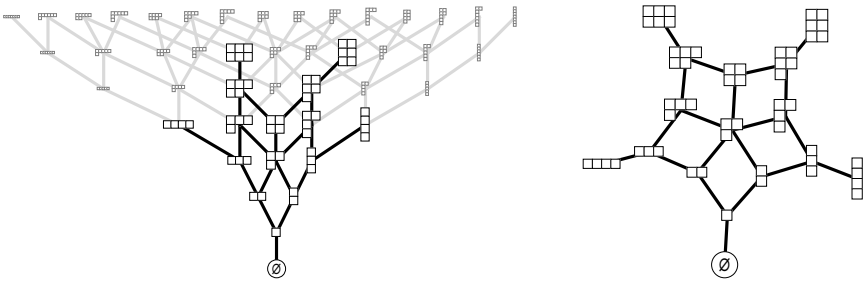

FIGURE 1. On the left is the subposet of Young's lattice contining those partitions with largest hook at most 4. On the right is the underlying graph, drawn to reveal a five-fold symmetry.

In a followup paper [Sut04], R. Suter explained this symmetry by relating this subposet to the weak order on the two-fold dilation of the fundamental alcove in affine type $\widetilde{A}_n$. The second author was introduced to this symmetry by V. Reiner and D. Stanton, and gave a combinatorial generalization in [Wil11] to accommodate a parameter m. The corresponding geometric interpretation was given by M. Zabrocki and C. Berg [BWZ11] as an m-fold dilation of the fundamental alcove. We refer the reader to [TW14, Section 2] for a more detailed history.

To characterize the orbit structure of this cyclic symmetry, the second author defined a map to words under rotation via a general technique called a *bijaction*. It isn't hard to show that this map is a bijection when the parameter m is equal to 2 [Wil11], and M. Visontai was able to invert it for $m = 3$ [VW12]. The general case was resistant to all attack—until the authors were introduced at the 2012

Combinatorial Algebra meets Algebraic Combinatorics conference at the Université du Québec à Montréal. Our joint paper [TW14] resulted.

Unexpectedly, it turns out that inverting the general sweep map and inverting the bijaction to describe the symmetry of the fundamental alcove in type $\widetilde{A}_n$ are essentially the same problem.

3. Presweeping and Its Inverse

We will factor the modular sweep map as the composition of two maps: the modular presweep map and the forgetful map. In this section, we define the modular presweep map and its inverse.

3.1. **Notation.** Adhering to the notation in [TW14], we prefer to think of the modular levels from the introduction as partitioning the word u into blocks. Define a *partitioned word* for $\mathsf{u} \in \mathcal{A}$ to be a partition u^* of u into m words $\mathsf{u}^* = \mathsf{u}^*_{m-1}|\mathsf{u}^*_{m-2}|\cdots|\mathsf{u}^*_0$—where we use the *block divider* symbol $|$ to separate the blocks—so that $\mathsf{u} = \mathsf{u}^*_{m-1}\cdots\mathsf{u}^*_0$ is their concatenation. We call the word u^*_k the kth *block* and, with apologies to the combinatorics of words community, we write $\mathcal{A}^*$ for the set of all partitioned words of $\mathcal{A}$. We may use either the symbol $\cdot$ or $\emptyset$ to denote an empty block. If the i-th letter u_i of u belongs to the k-th block u^*_k in the partitioned word u^*, we let $\mathsf{block}(\mathsf{u}^*, i) := k$. We fix the notation $|\mathsf{u}| := \sum_{i=1}^N \mathsf{u}_i$ and $|\mathsf{u}|_m = \ell_N = |\mathsf{u}| \bmod m$.

3.2. **The Modular Presweep Map.** The *modular presweep map* is the function $\mathsf{presweep}_m : \mathcal{A} \hookrightarrow \mathcal{A}^*$ that sorts $\mathsf{w} \in \mathcal{A}$ into blocks according to its levels. Precisely, for $k = m-1, \ldots, 2, 1, 0$, first initialize $\mathsf{u}^* := \cdot|\cdot|\cdots|\cdot$ to be the empty partitioned word, then read w from right to left and append to u^*_k all letters w_j whose level $\ell_j = \left(\sum_{i=1}^j \mathsf{w}_i \bmod m\right)$ is equal to k. In other words, u^*_k is obtained by extracting all letters of level k from u and reversing their relative order. Pseudo-code for $\mathsf{presweep}_m$ is given in Algorithm 1.

Input: A word $\mathsf{w} = \mathsf{w}_1\mathsf{w}_2\cdots\mathsf{w}_N \in \mathcal{A}$.
Output: A partitioned word $\mathsf{u}^* = \mathsf{u}^*_{m-1}|\mathsf{u}^*_{m-2}|\cdots|\mathsf{u}^*_0 \in \mathcal{A}^*$.
Let $\mathsf{u}^* := \cdot|\cdot|\cdots|\cdot$;
for $k = m-1$ **to** 0 **do**
  for $j = N$ **to** 1 **do**
    if $\left(\sum_{i=1}^j \mathsf{w}_i \bmod m\right) = k$ **then**
      Append w_j to u^*_k;
    end
  end
end
Return u^*;

Algorithm 1: $\mathsf{presweep}_m : \mathcal{A} \hookrightarrow \mathcal{A}^*$.

Example 3.1. As in Example 1.3, let $m = 5$, $N = 7$, $\mathsf{w} = 3113214 \in \mathcal{A}$, and we compute the modular levels of a word $\mathsf{w} \in A$ by summing the initial letters of w modulo m:

w	3	1	1	3	2	1	4
ℓ	3	4	0	3	0	1	0

.

We now compute the modular sweep of w by sorting by levels, reading w from right to left. Placing letters with the same level in a block, we obtain the corresponding partitioned word $\mathsf{u}^* := \mathsf{presweep}_m(\mathsf{w})$ in $\mathcal{A}^*$:

$$\begin{array}{c|c|cc|c|c|ccc} \ell & 4 & 3 & 3 & 2 & 1 & 0 & 0 & 0 \\ \mathsf{u}^* & 1 & 3 & 3 & \cdot & 1 & 4 & 2 & 1 \end{array}.$$

3.3. **The Inverse Modular Presweep Map.** The *inverse modular presweep map* is the function $\mathsf{inverse_presweep}_m : \mathcal{A}^* \to \mathcal{A}$ such that

$$\mathsf{inverse_presweep}_m \circ \mathsf{presweep}_m = \mathsf{id}_{\mathcal{A}},$$

where $\mathsf{id}_{\mathcal{A}}(\mathsf{w}) = \mathsf{w}$ is the identity function on $\mathcal{A}$. As explained in [ALW14b, Section 5.2] and in [TW14, Algorithm 2 and Figure 8], if we know how to associate the correct levels to $\mathsf{u} := \mathsf{sweep}_m(\mathsf{w})$, it is easy to reconstruct w.

Suppose we have the partitioned word $\mathsf{u}^* := \mathsf{presweep}_m(\mathsf{w})$. Since the letters of w were just rearranged to make u^*, we can determine $\ell_N = |\mathsf{w}|_m = |\mathsf{u}|_m$ from u^*. As we swept w from right to left, the last letter of w is therefore the first letter in $\mathsf{u}^*_{\ell_N}$. Remove this letter from u^*. Subtracting this letter from ℓ_N gives ℓ_{N-1}, and we obtain the $(N-1)$st letter of w as the first remaining letter in block ℓ_{N-1}. In general, for $i = 1, 2, \ldots, N$ we have already computed ℓ_{N-i+1}; subtracting the leftmost remaining letter in $\mathsf{u}^*_{\ell_{N-i+1}}$ from ℓ_{N-i+1} (and removing it from $\mathsf{u}^*_{\ell_{N-i+1}}$) gives ℓ_{N-i}. Pseudo-code for $\mathsf{inverse_presweep}_m$ is given in Algorithm 2.

Input: A partitioned word $\mathsf{u}^* = \mathsf{u}^*_{m-1}|\mathsf{u}^*_{m-2}|\cdots|\mathsf{u}^*_0 \in \mathcal{A}^*$.
Output: A word $\mathsf{w} = \mathsf{w}_1\mathsf{w}_2\cdots\mathsf{w}_N \in \mathcal{A}$ or a subword of u^*.
Let $\ell_N := \left(\sum_{j=1}^{N} \mathsf{u}_j \bmod m\right)$ and $\mathsf{w} := \emptyset$;
for $i = 1$ **to** N **do**
  if $\mathsf{u}^*_{\ell_{N-i+1}} \neq \emptyset$ **then**
    Remove the first letter of $\mathsf{u}^*_{\ell_{N-i+1}}$ and assign it to w_{N-i+1};
    Prepend w_{N-i+1} to w;
    Let $\ell_{N-i} := (\ell_{N-i+1} - \mathsf{w}_{N-i+1} \bmod m)$;
  end
  else
    Return u^*
  end
end
Return w;

Algorithm 2: $\mathsf{inverse_presweep}_m : \mathcal{A}^* \hookrightarrow \mathcal{A}$.

We say that Algorithm 2 *succeeds* on a partitioned word u^* if it returns an element of $\mathcal{A}$, and we say that it *fails* if it returns an element of $\mathcal{A}^*$. We call a partitioned word u^* on which Algorithm 2 succeeds a *successful partition* of the underlying word u. Since Algorithm 2 undoes Algorithm 1 one step at a time, we conclude that $\mathsf{presweep}_m$ and $\mathsf{inverse_presweep}_m$ are inverses.

Example 3.2. To reverse Example 3.1, we compute as follows. We start with the partitioned word u^*:

$$\begin{array}{c|c|cc|c|c|ccc} \ell & 4 & 3 & 3 & 2 & 1 & 0 & 0 & 0 \\ \mathsf{u}^* & 1 & 3 & 3 & \cdot & 1 & 4 & 2 & 1 \end{array}.$$

We find $\ell_N = |\mathsf{u}|_m = 0$. Then we iterate:

i	u^*	ℓ_{N-i}	w
0	$1\|33\| \cdot \|1\|\dot{4}21$	0	$\cdot$
1	$1\|33\| \cdot \|\dot{1}\|421$	1	4
2	$1\|33\| \cdot \|1\|4\dot{2}1$	0	14
3	$1\|\dot{3}3\| \cdot \|1\|421$	3	214
4	$1\|33\| \cdot \|1\|42\dot{1}$	0	3214
5	$\dot{1}\|33\| \cdot \|1\|421$	4	13214
6	$1\|3\dot{3}\| \cdot \|1\|421$	3	113214
7	$1\|33\| \cdot \|1\|421$	0	3113214

Comparing with Example 3.1, we see that we have recovered w.

3.4. **Forgetting.** We now obtain the modular sweep map from the modular presweep map by forgetting the information of the blocks. The *forgetful map* is the function

$$\mathsf{forget} : \mathcal{A}^* \to \mathcal{A}$$
$$\mathsf{forget}\left(\mathsf{u}^*_{m-1}|\mathsf{u}^*_{m-2}|\cdots|\mathsf{u}^*_0\right) = \mathsf{u}^*_{m-1}\mathsf{u}^*_{m-2}\cdots\mathsf{u}^*_0$$

obtained by concatenating all the blocks of $\mathsf{u}^* \in \mathcal{A}^*$. Thus, the modular sweep map of Section 1 may be written as the composition

$$\mathsf{sweep}_m = (\mathsf{forget} \circ \mathsf{presweep}_m) : \mathcal{A} \to \mathcal{A}.$$

Example 3.3. Continuing with Example 3.1, we forget the partitioning to obtain the modular sweep $\mathsf{u} := \mathsf{sweep}_m(\mathsf{w})$ of w to be

$$(\mathsf{forget} \circ \mathsf{presweep}_m)(3113124) = \mathsf{forget}\,(1|33| \cdot |1|421) = 1331421.$$

Thus, the problem of inverting the modular sweep map has been reduced to showing that there exists a unique successful partition $\mathsf{u}^* \in \mathcal{A}^*$ for each word $\mathsf{u} \in \mathcal{A}$. We do this in the next section.

4. Equitable Partitions and the Successful Partition

We already solved the problem of constructing the successful partition in [TW14], where we studied a composition

$$f \circ p,$$

where f is the map forget and p is a map *very slightly* different from $\mathsf{presweep}_m$ (see Appendix A). In particular, our notions here of a successful partition and the forgetful map coincide with those in [TW14].

Our algorithm for inverting forget passes by way of the notion of *equitable partitions*—starting with any equitable partition as an initial approximation to the inverse, Algorithm 4 modifies this guess until it converges to the actual preimage (the successful partition of Section 3.3). As we prove in Theorem 4.12, it happens that our initial approximation (the rightmost partition, constructed in Algorithm 3 using an idea of F. Aigner, C. Ceballos, and R. Sulzgruber) is already the desired preimage.

4.1. **Equitable Partitions.** We expand a partitioned word u^* into an $N \times m$ *balancing array*

$$M^{\mathsf{u}^*} = (M^{\mathsf{u}^*}_{i,j})_{\substack{1 \le i \le N \\ m-1 \ge j \ge 0}}$$

defined by

$$M^{\mathsf{u}^*}_{i,j} := \begin{cases} \blacksquare & \text{if } j \in \{\mathsf{block}(\mathsf{u}^*, i), \mathsf{block}(\mathsf{u}^*, i)+1, \dots, \mathsf{block}(\mathsf{u}^*, i)+\mathsf{u}_i - 1\} \bmod m, \\ \cdot & \text{otherwise,} \end{cases}$$

We say that column j (for $m-1 \ge j \ge 0$) of M^{u^*} is *equitably filled* if:

- $|\mathsf{u}|_m \ge j \ge 1$ and column j has $\left\lfloor \frac{|\mathsf{u}|}{m} \right\rfloor + 1$ copies of the symbol $\blacksquare$, or if
- $j = 0$ or $j > |\mathsf{u}|_m$, and column j contains $\left\lfloor \frac{|\mathsf{u}|}{m} \right\rfloor$ copies of $\blacksquare$.

If a column has (strictly) fewer copies of the symbol $\blacksquare$ than it would to be equitably filled, we say it is *less than equitably filled*; similarly, when a column has (strictly) more copies of $\blacksquare$ we say that it is *more than equitably filled*. In particular, if $|\mathsf{u}|_m = 0$, then every column has $\left\lfloor \frac{|\mathsf{u}|}{m} \right\rfloor$ copies of $\blacksquare$. We say that u^* is an *equitable partition* if each of the columns of M^{u^*} is equitably filled.

The motivation for this definition is the following lemma.

Lemma 4.1. *Any successful partition u^* is an equitable partition.*

Proof. We can construct all successful partitions as follows [TW14, Definition 7.5]. Define an infinite complete m-ary tree $\mathcal{T}^*_m$ by

(1) The zeroth rank consists of the empty successful partition u^*, given by $\mathsf{u}^*_k = \emptyset$ for $m-1 \ge k \ge 0$.
(2) The children of a successful partition $\mathsf{u}^* = \mathsf{u}^*_{m-1}|\dots|\mathsf{u}^*_0$ are the m successful partitions obtained by prepending $i \pmod m$ to $\mathsf{u}^*_{i+|\mathsf{u}|_m}$.

Then it is easy to see that all partitioned words in $\mathcal{T}^*_m$ are equitable, and that all successful partitions appear in $\mathcal{T}^*_m$ [TW14, Lemma 7.2]. (It is not yet clear that the images of the words in $\mathcal{T}^*_m$ under the forgetful map are actually distinct.) □

Example 4.2. As in Example 1.1, the equitable partitions $13|31|4|2|1$ and $1|33|\cdot|1|421$ have corresponding balancing arrays

		j				
		4	3	2	1	0
	1	■	·	·	·	·
	2	■	■	■	·	·
	3	·	■	■	■	·
i	4	·	■	·	·	·
	5	■	·	■	■	■
	6	·	·	·	■	■
	7	·	·	·	·	■

and

		j				
		4	3	2	1	0
	1	■	·	·	·	·
	2	·	■	■	■	·
	3	·	■	■	■	·
i	4	·	·	·	■	·
	5	■	■	■	·	■
	6	■	·	·	·	■
	7	·	·	·	·	■

.

Since $\left\lfloor \frac{|\mathsf{u}|}{m} \right\rfloor = \left\lfloor \frac{15}{5} \right\rfloor = 3$ and $|\mathsf{u}|_m = 0$, any equitable filling has three copies of $\blacksquare$ in each column j. Note that the labeling of columns has changed from Example 1.1.

It is clear that the notion of equitable partitions generalizes the "real world" example given in Section 1.1 (the difference being that we no longer assume that $|\mathsf{u}|$ is divisible by m).

4.2. **The Rightmost Partition.** Given u, we first prove the existence of a particular equitable partition.

Definition 4.3. A *rightmost equitable partition* is an equitable partition u^* such that any other equitable partition v^* has $\mathsf{block}(\mathsf{v}^*, i) \geq \mathsf{block}(\mathsf{u}^*, i)$ for all i.

Theorem 4.4. *Any* $\mathsf{u} \in \mathcal{A}$ *admits a unique rightmost partition* $\mathsf{rightmost}(\mathsf{u})$.

Proof. We claim that Algorithm 3 constructs the unique rightmost equitable partition.[3]

Input: A word $\mathsf{u} \in \mathcal{A}$.
Output: The rightmost equitable partition $\mathsf{u}^* \in \mathcal{A}^*$.
Set $\mathsf{u}^* = \cdot|\cdot|\cdots|\mathsf{u}$;
while u^* *is not an equitable partition* **do**
  Let j be the rightmost column of M^{u^*} that is less than equitably filled;
  Delete the leftmost letter of u^*_{j-1} and append it to u^*_j;
end
Return u^*;

Algorithm 3: $\mathsf{rightmost} : \mathcal{A} \to \mathcal{A}^*$.

We begin with a lemma about which columns can be more than equitably filled during the execution of Algorithm 3.

Lemma 4.5. *While running Algorithm 3, before we have reached an equitable filling, column 0 will always be more than equitably filled.*

Proof. The set of columns which are more than equitably filled can only decrease over time, since we never make a move which increases the number of copies of ■ in a column above the equitable amount. Thus, if a column is more than equitably filled at a certain point, it has been more than equitably filled since the beginning.

Consider some $j > 0$, and suppose that at some point in the execution of Algorithm 3, column j is more than equitably filled. By what we have just argued, column j has been more than equitably filled from the beginning. Thus, we never pushed any letters from the $(j-1)$st block into the jth block, and we never moved any letters into any block to the left of it (since any such letter would have had to pass through the jth block). So the reason that the jth column is more than equitably filled is accounted for entirely by wrap-around from letters which are in blocks to its right. Every column to the left of j, and column 0, will also receive copies of ■ from each of the letters which contribute to the jth column. It follows that they are all more than equitably filled.

Thus, at each step of the algorithm, if the filling is not yet equitable, the columns which are more than equitably filled consist of an initial (left-aligned) subsequence of the columns, possibly empty, together with column 0. □

With this lemma, we first show that Algorithm 3 does not attempt any illegal moves, and so returns an equitable partition. Let j be the rightmost column which is less than equitably filled. Algorithm 3 might fail in two ways:

(1) Suppose $j = 0$. If column 0 is less than equitably filled, then Algorithm 3 would attempt to delete the leftmost letter of the nonexistent block u^*_{-1}. But this case cannot arise by Lemma 4.5.

[3]Algorithm 3 was communicated to us by F. Aigner, C. Ceballos, and R. Sulzgruber; although we had suspected the existence of Algorithm 3 (roughly dual to our Algorithm 5) from our analogy with stable marriages (see Remark 5.10), we had been unable to produce it.

(2) Now suppose $j > 0$. Algorithm 3 might try to delete the leftmost entry of an *empty* block u^*_{j-1} when column $j > 0$ is less than equitably filled.

If $j > 1$, since column $j-1$ is to the right of j, it must be at least equitably filled. Since the number of copies of ■ in an equitable filling of column j at most the number in an equitable filling of column $j-1$, and column $j-1$ is equitably filled while j is not, there must be a letter in block $j-1$.

If $j = 1$, the above argument fails because the number of copies of ■ in an equitable filling of column 1 may be greater (by one) than the number in an equitable filling of column 0. But by Lemma 4.5, we know that column 0 is strictly more than equitably filled, and so if column 1 is less than equitably filled, there must be a letter in block 0.

We now show that Algorithm 3 outputs the unique rightmost equitable partition u^*. For each step of Algorithm 3, record the pair (j, i), where j is the rightmost column that is less than equitably filled, and i is the index of the letter u_i that we move from u^*_{j-1} to u^*_j. Consider another equitable partition v^* of u and find the first recorded pair (j, i) such that u_i is in v^*_{j-1}. If there is no such pair, then each letter of v^* is in a block at least as far to the left as in u^* and we are done.

Otherwise, write u^∂ for the partition produced by running Algorithm 3 up to (but not including) the step that produced the pair (j, i). The partition u^∂ is non-equitable, and j is its rightmost column that is less than equitably filled. Every letter in a block to the left of j that does not contribute a ■ to column j for u^∂, is, by our choice of pair (j, i), at least as far to the left in v^* as in u^∂—and so does not contribute a ■ to column j in v^* either. Every letter in a block to the right of j is, again by our choice of (j, i), also at least as far to the left. At the same time, since the initial letter in $\mathsf{u}^\partial_{j-1}$ is by assumption in v^*_{j-1}, no letter that does not contribute a ■ to column j for u^∂ can have moved to the left of v^*_{j-1} to contribute a ■ to column v^*_j. Thus, v^* has at most as many copies of ■ in column j as u^∂ does, which is too few to be equitable, contradicting our assumption that v^* was equitable.

□

Example 4.6. We illustrate Algorithm 3 applied to the word $\mathsf{u} = 1331421$. At each step, the rightmost column with less than its equitable filling is highlighted.

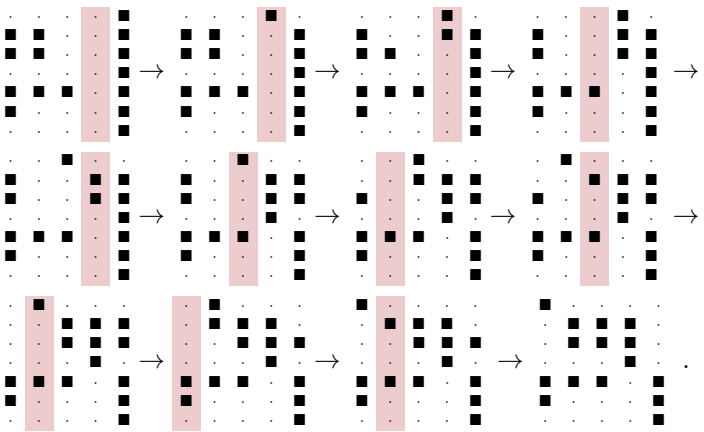

Thus, the rightmost equitable partition of u is $\mathsf{u}^* = 1|33| \cdot |1|421$.

4.3. **Balanced Block Suffixes.** In this section, we give some preparatory structural definitions and results on equitable partitions.

Definition 4.7. Fix $\mathsf{u} \in \mathcal{A}$ and an equitable partition u^* of u. A *balanced block-suffix* is a partitioned subword s^* of u^* with $|\mathsf{s}^*|_m = 0$, whose intersection with each block of u^* is a suffix of that block. We can view s^* as an equitable partition of a (not-necessarily consecutive) subword s of u.

A *left balanced block-suffix* has the additional constraint that it has empty rightmost block: $\mathsf{s}_0^* = \emptyset$. A *minimal balanced block-suffix* is a balanced block-suffix whose intersection with all other balanced block-suffixes is itself or empty.

Note that the condition $|\mathsf{s}^*|_m = 0$ implies that balanced block-suffixes contribute the same number of copies of ■ to each column of M_{s^*}.

Example 4.8. The equitable partition 13|31|4|2|1 has three balanced block-suffixes:

$$\begin{array}{rcl} \mathsf{s}^* & = & 1\ 3 \mid 3\ 1 \mid 4 \mid 2 \mid 1 \\ \mathsf{t}^* & = & 1\ 3 \mid 3\ 1 \mid 4 \mid 2 \mid 1 \\ \mathsf{s}^* \cup \mathsf{t}^* & = & 1\ 3 \mid 3\ 1 \mid 4 \mid 2 \mid 1 \end{array}.$$

All three are left balanced block-suffixes, but only the first two are minimal.

The following lemma establishes the existence of a *maximal left balanced block-suffix*—that is, a left balanced block suffix containing all left balanced block suffixes.

Lemma 4.9. *Let u^* be an equitable partition that is not successful. Then the output of Algorithm 2 is the maximal left balanced block-suffix s^* of u^*.*

Proof. Algorithm 2 terminates when it tries to remove a letter from an empty block of u^*. This block corresponds to a column in M_{u^*} whose copies of ■ all came from other parts. Since Algorithm 2 preserves the property of being equitable, it must remove a letter from a block corresponding to the leftmost column with the most copies of ■, or a letter from a block corresponding to the rightmost column (if all columns have the same number of copies of ■). Then Algorithm 2 does not succeed only if it was trying to remove a letter from the empty rightmost block, so that the copies of ■ corresponding to the remaining letters in u^* are equally distributed among the columns, and there are none in the rightmost block. This is the condition to be a left balanced block-suffix.

Moreover, since any left balanced block-suffix s^* has no element in its rightmost block s_0^*, it will be untouched by Algorithm 2—since s^* is equitable, $|\mathsf{s}|_m = 0$, and s^* is a suffix of each block of u^*, removing s^* from u^* does not affect how Algorithm 2 runs. Therefore, the output of Algorithm 2 is the maximal left balanced block-suffix. □

4.4. **The Successful Partition.** In this section, we construct a successful partition, and show that it is unique. Our method is to start with any equitable partition as an approximation to the successful partition and apply Algorithm 2 to converge to the successful partition. In fact, the rightmost equitable partition is the successful partition.

Theorem 4.10 ([TW14, Proposition 7.3])**.** *Any $\mathsf{u} \in \mathcal{A}$ admits a unique successful partition* $\mathsf{successful}(\mathsf{u})$.

Proof. We claim that Algorithm 4 constructs a successful partition when any equitable partition is given as input (equitable partitions always exist, since we have constructed an example of one in the previous section). In words, Algorithm 4 proceeds as follows. By Lemma 4.1, we know that every successful partition is equitable. Beginning with u^* equal to the leftmost equitable partition of u, we run through Algorithm 2 "as far as possible." If Algorithm 2 does not succeed, we leave in place all letters of u^* that were visited, and shift all other letters of u^* to the

Input: A word $\mathsf{u} \in \mathcal{A}$.
Output: A successful partition $\mathsf{u}^* \in \mathcal{A}^*$.
Let u^* be any equitable partition of u;
while u^* *is not a successful partition* **do**
  Let s^* be the output from Algorithm 2 applied to u^*;
  for $t = m - 1$ **to** 0 **do**
    Delete the rightmost $|\mathsf{s}_t^*|$ entries from u_t^*;
    Prepend s_t^* to u_{t-1}^*;
  end
end
Return u^*;

Algorithm 4: $\mathsf{successful} : \mathcal{A} \to \mathcal{A}$.

right by one block. We shall show that this process is both finite and cannot fail, and so therefore eventually terminates in a successful partition.

By construction—assuming it is well-defined—Algorithm 4 returns a successful partition for any word $\mathsf{u} \in \mathcal{A}$. Let s^* be the partitioned word containing those letters that remain in u^* after applying Algorithm 2. By Lemma 4.9, these form a left balanced block-suffix. Shifting the letters in s^* one block to the right—while preserving those letters in u^* that do not appear in s^*—we obtain a new equitable partition for u. Each time we repeat this process, the letters are moved further to the right, so we never obtain the same equitable partition twice. Since there are only a finite number of partitioned words, Algorithm 4 must eventually terminate in a successful partition u^*.

We conclude uniqueness using the cardinality argument from [TW14]; a separate, direct argument is given below in Remark 5.9. By Algorithm 4, every word in $\mathcal{A}$ has a successful partition. Since the tree $\mathcal{T}_m^*$ in the proof of Lemma 4.1 contains every successful partition of words in $\mathcal{A}$, and since its Nth level has m^N elements, we conclude that every word in $\mathcal{A}$ has a unique successful partition. □

Example 4.11. We compute the successful partition starting from the equitable partition 13|31|4|2|1. The rightmost column indicates which letters were visited by Algorithm 2, and in what order (the corresponding rows are colored gray). The key to the proof that Algorithm 4 is well-defined is that whenever Algorithm 2 finishes, the copies of ■ corresponding to the *remaining* letters (which are coloured black in the diagram) are equally distributed among the columns.

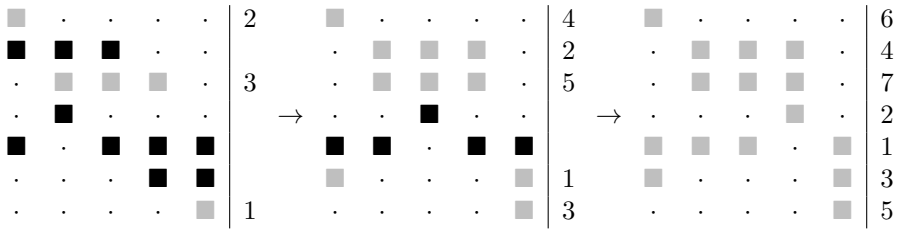

Thus, the successful partition of u is $\mathsf{successful}(\mathsf{u}) = 1|33| \cdot |1|421$—which agrees with Examples 3.1 and 3.2.

Theorem 4.12. *For* $\mathsf{u} \in \mathcal{A}$, $\mathsf{rightmost}(\mathsf{u}) = \mathsf{successful}(\mathsf{u})$.

Proof. If an equitable partition were rightmost without being successful, shifting the letters not visited by Algorithm 2 over gives a new equitable partition further to the right—which is a contradiction. □

4.5. **The Modular Sweep Map.** We can now prove our main theorem.

Theorem 1.4. *The modular sweep map is a bijection $\mathcal{A} \to \mathcal{A}$.*

Proof. Recall that the modular sweep map may be written as the composition

$$\mathsf{sweep}_m = \mathsf{forget} \circ \mathsf{presweep}_m : \mathcal{A} \to \mathcal{A}.$$

Since Section 3.3 inverts $\mathsf{presweep}_m$ and Theorem 4.10 inverts forget, we conclude that the modular sweep map is invertible and is given by the composition of the remembering map with the inverse of the modular presweep map. □

5. The Distributive Lattice of Equitable Partitions

In this section, motivated by an analogy to stable marriages, we endowing the set of equitable partitions with the structure of a distributive lattice.

5.1. **The Leftmost Equitable Partition.** Having defined the *rightmost* equitable partition, it is natural to ask for the *leftmost* equitable partition. As it turns out, there is exactly one leftmost equitable partition.

Definition 5.1 ([TW14, Definition 8.3])**.** A *leftmost equitable partition* is an equitable partition u^* such that any other equitable partition v^* has $\mathsf{block}(\mathsf{v}^*, i) \leq \mathsf{block}(\mathsf{u}^*, i)$ for all i.

Theorem 5.2 ([TW14, Lemma 8.1])**.** *For any word $\mathsf{u} \in A$, there is a unique leftmost equitable partition $\mathsf{leftmost}(\mathsf{u})$.*

Proof. We claim that Algorithm 5 constructs the unique leftmost equitable partition.

Input: A word $\mathsf{u} \in \mathcal{A}$.
Output: The leftmost equitable partition $\mathsf{u}^* \in \mathcal{A}^*$.
Set $\mathsf{u}^* = \mathsf{u}|\cdot|\cdot|\cdots|\cdot$;
while u^* *is not an equitable partition* **do**
  Let j be the leftmost column of M^{u^*} that is more than equitably filled;
  Delete the rightmost letter of u^*_j and prepend it to u^*_{j-1};
end
Return u^*;

Algorithm 5: $\mathsf{leftmost} : \mathcal{A} \to \mathcal{A}^*$.

Lemma 5.3. *While running Algorithm 5, column 0 will always be less than equitably filled or equitably filled.*

Proof. The set of columns which are less than equitably filled or equitably filled can only decrease over time, since we never make a move which decreases the number of copies of ■ in a column below the equitable amount. Thus, if a column is less than equitably filled at a certain point, it has been less than equitably filled since the beginning.

Consider some $j > 0$, and suppose that at some point in the execution of Algorithm 5, column j is less than equitably filled. By what we have just argued, column j has been less than equitably filled from the beginning. In particular, we have never pushed any letters from the jth block to the $(j-1)$st block, and we have never moved any letters into any block to the right of it (since any such letter would have had to pass through the jth block). So every column to the right of column j, including column 0, contains at most the same number of copies of ■ as column j. It follows that all columns to the right of column j—except possibly

column 0—are all also less than equitably filled. Since an equitable filling of column j has more than or equal to an equitable filling of column 0, column 0 is either less than equitably filled or equitably filled.

Thus, at each step of the algorithm, if the filling is not yet equitable, column 0 is either less than equitably filled or equitably filled. □

With this lemma, we first show that Algorithm 5 does not attempt any illegal moves, and so returns an equitable partition. Let j be the leftmost column that is more than equitably filled. Algorithm 5 might fail in two ways:

(1) Suppose $j = 0$. Then Algorithm 5 would attempt to delete the rightmost letter of u_0^* and append it to the nonexistent block u_{-1}^*. But this case cannot arise by Lemma 5.3.
(2) Otherwise, suppose $j > 0$. Algorithm 5 might try to delete the rightmost entry of an *empty* block u_j^* when column j is more than equitably filled. Since u_j^* is empty, the number of copies of ■ in column j is at most the number in column $j+1$. Therefore, this case arises only when an equitable filling of column j has strictly fewer copies of ■ than an equitable filling of column $j+1$. But the only column with this property is column 0.

We now show that Algorithm 5 outputs the unique leftmost equitable partition u^*. For each step of Algorithm 5, record the pair (j, i), where j is the leftmost column that is more than equitably filled, and i is the index of the letter u_i that we move from u_j^* to u_{j-1}^*. Consider another equitable partition v^* of u and find the first recorded pair (j, i) such that u_i is in v_j^*. If there is no such pair, then each letter of v^* is in a block at least as far to the right as in u^* and we are done.

Otherwise, write u^∂ for the partition produced by running Algorithm 5 up to (but not including) the step that produced the pair (j, i). The partition u^∂ is non-equitable, and j is its leftmost column that is more than equitably filled. Every letter that contributes a ■ to column j for u^∂, is, by our choice of pair (j, i), at least as far to the right in v^* as in u^∂. At the same time, since the final letter in u_j^∂ is by assumption also in v_j^*, none of the letters that contribute a ■ to column j for u^∂ can have moved to the right of v_j^*. Thus, v^* has at least as many copies of ■ in column j as u^∂ does, which is too many to be equitable, contradicting our assumption that v^* was equitable. □

Example 5.4. We illustrate Algorithm 5 applied to the word $\mathsf{u} = 1331421$. At each step, the leftmost column with more than its equitable filling is highlighted.

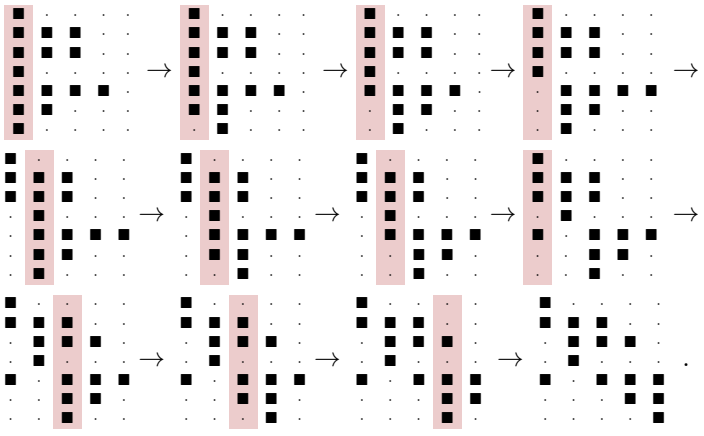

Thus, the leftmost equitable partition of u is $\mathsf{u}^* = 13|31|4|2|1$.

5.2. **The Distributive Lattice and Stable Marriages.** Fix a word u. We define the *componentwise order* on the set of equitable partitioned words of u by $\mathsf{u}^* \leq \mathsf{v}^*$ iff $\mathsf{block}(\mathsf{u}^*, i) \leq \mathsf{block}(\mathsf{v}^*, i)$ for all $1 \leq i \leq N$.

Example 5.5. The componentwise order on the five equitable partitions for the word $\mathsf{u} = 1331421$ appears below. The rightmost equitable partition computed in Example 4.6 is the maximal element, while the leftmost equitable partition computed in Example 5.4 is the minimal element. For each equitable partition, the letters read by Algorithm 2 are decorated with a dot. Algorithm 4—run in Example 4.11—visits those equitable partitions highlighted in red. Letters of a minimal left balanced block-suffixes are given an underline (or wavy underline). Observe the the cover relations are given by shifting minimal left balanced block-suffixes.

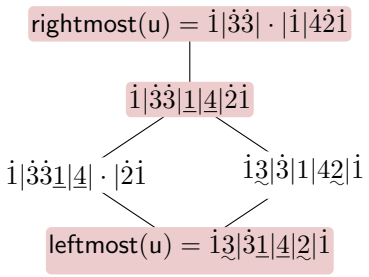

Lemma 5.6. *Let s^* be a minimal left balanced block-suffix of u^*, and let v^* be the result of shifting the letters in s^* one block to the the right. Then $\mathsf{u}^* \lessdot \mathsf{v}^*$ is a cover relation in the componentwise order. All cover relations are of this form.*

Proof. The first statement follows from the definition of a minimal left balanced block-suffix—once we move any member of that block-suffix, we are forced to move them all to preserve equitability.

To establish the converse, suppose that we have a cover relation $\mathsf{u}^* \lessdot \mathsf{v}^*$. Let s^* be the letters of u^* which have moved to the right in v^*. We want to establish that each of the letters of s^* has moved exactly one step to the right in v^*. Choose some letter u_i of s^* rightmost in its block, and consider the cascade of forced moves that follow once we move it one block to the right: specifically, moving it to the right means that some column now has too many copies of ■, which forces a letter to move out of that block, which means that another column now has too many copies of ■, and so on. We refer to this sequence as the cascade of forced moves. The cascade never tries to move a letter out of an empty block because there is necessarily an element of s^* available whenever it is needed. It therefore terminates at an equitable partition, which must be v^*, by our assumption that $\mathsf{u}^* \lessdot \mathsf{v}^*$.

Suppose that up to a certain step in the cascade, we have not shifted any letter more than once, but we now move a letter u_j for the second time. At this point, let u^∂ be the partition and let $\mathsf{block}(\mathsf{u}^\partial, j) = k$. Suppose that $|\mathsf{u}^*_k| = \ell$, so that all ℓ letters originally in block u^*_k have already been shifted. Then the block to which our first letter u_i belonged—$\mathsf{block}(\mathsf{u}^*, i)$—is not equal to k, because the cascade would have concluded upon arriving back in block $\mathsf{block}(\mathsf{u}^*, i)$. Also, since it is impossible to shift letters from the rightmost block, we must have $k \neq 0$. In order to shift a letter from block k, we must have just finished shifting a letter whose corresponding row in M^{u^∂} has rightmost ■ in column $k-1$. Since u^* is equitable, we know that in u^*, the number of letters whose rightmost ■ is in column $k-1$ is equal to or one less than the number of letters in u^*_k. By our assumption that we have moved letters at most once in our cascade from u^* to u^∂, it follows that we have previously moved at most $\ell - 1$ letters out of block u^*_k—one fewer than the number which it

started with, contrary to our assumption that at this time, we had already moved all of them.

Since each of the letters of s^* is moved exactly once, it must be a left balanced block-suffix. It is then clear that it must also be minimal. □

Given two equitable partitions u^* and v^*, let $\mathsf{u}^* \vee \mathsf{v}^*$ be the partition given by

$$\mathsf{block}(\mathsf{u}^* \vee \mathsf{v}^*, i) := \min\{\mathsf{block}(\mathsf{u}^*, i), \mathsf{block}(\mathsf{v}^*, i)\},$$

and let $\mathsf{u}^* \wedge \mathsf{v}^*$ be the partition given by

$$\mathsf{block}(\mathsf{u}^* \wedge \mathsf{v}^*, i) := \max\{\mathsf{block}(\mathsf{u}^*, i), \mathsf{block}(\mathsf{v}^*, i)\}.$$

Proposition 5.7. *For two equitable partitions* u^* *and* v^**, both* $\mathsf{u}^* \vee \mathsf{v}^*$ *and* $\mathsf{u}^* \wedge \mathsf{v}^*$ *are equitable partitions.*

Proof. We give the argument for $\mathsf{u}^* \wedge \mathsf{v}^*$; the dual argument gives the corresponding result for $\mathsf{u}^* \vee \mathsf{v}^*$. The leftmost equitable partition $\mathsf{leftmost}(\mathsf{u})$ is certainly a lower bound for both u^* and v^*. Suppose there exists an element u_i at the end of a block in $\mathsf{leftmost}(\mathsf{u})$ such that $\mathsf{block}(\mathsf{leftmost}(\mathsf{u}), i) \neq \max\{\mathsf{block}(\mathsf{u}^*, i), \mathsf{block}(\mathsf{v}^*, i)\}$. Then u_i is a member of a left balanced block-suffix s^* of $\mathsf{leftmost}(\mathsf{u})$, and we may move all of the letters in s^* one block to the right to move up in the componentwise order. The resulting partition is still a lower bound for both u^* and v^*, and we may continue this process so long as there is an element u_i such that $\mathsf{block}(\mathsf{leftmost}(\mathsf{u}), i) \neq \max\{\mathsf{block}(\mathsf{u}^*, i), \mathsf{block}(\mathsf{v}^*, i)\}$. The process therefore terminates with $\mathsf{u}^* \wedge \mathsf{v}^*$. □

We conclude the following.

Theorem 5.8. *The set of equitable partitions of a fixed word* u *forms a distributive lattice under componentwise order. The minimal element is the leftmost equitable partition and the maximal element is the rightmost equitable partition.*

Proof. Since the operations max and min distribute, we obtain the first statement. The second statement follows from Theorem 4.4 and Theorem 5.2. □

Remark 5.9. It is now possible to give a direct argument of the uniqueness of the successful partition as follows. Let u^* be a partition of u which is not rightmost. It therefore admits a cover in the lattice of equitable partitions. By Lemma 5.6, this means that it has a left balanced block suffix, which implies that it is not successful.

Remark 5.10. Theorem 5.8 was found by analogy with stable marriages, as we now explain. The theory of stable marriages is due to D. Gale and L. Shapley [GS62]. Briefly, given n men and n women who have each individually totally ordered the opposite sex, an *unstable set of marriages* is a bijection between the men and the women such that there exist a man and a woman (not married to each other) who prefer each other to their actual spouses. A *stable set of marriages* is a bijection between the men and women which is not unstable. In [GS62], D. Gale and L. Shapley proved the existence of a set of stable marriages by giving an algorithm that produced one. This work of D. Gale and L. Shapley has since found diverse applications—for example, it provided the mathematical underpinning for reform of residency matches, such as the San Francisco match for ophthalmology and the National Resident Matching Program (NRMP), which assign graduating medical students to residency programs.

Although the men and women appear to be equals in the statement of the problem, D. Gale and L. Shapley's algorithm has the interesting property of not treating the two groups symmetrically. In their algorithm, the members of one group act by proposing to the members of the other group—working down from the top of their list—while the other group continously rejects all but the most favorable proposal

they have yet received. This asymmetry is highlighted by the outcome: when the men propose to the women, the output is the worst possible set of stable marriages for the women (in the sense that in any other possible set of stable marriages, every woman would be paired with a man no lower in her total order) and the best possible for the men. This bias was originally raised in reference to the NRMP in [WWW81], where it was shown that the choice had been made to favor the interests of residency programs over those of students. This bias has now been largely switched, with the consent of all parties; the San Francisco match switched in 1996 [Wil96a, Wil96b], while the larger NRMP switched under public pressure two years later.

In [Knu76], D. Knuth credits J. Conway with the observation that the set of all stable sets of marriages forms a distributive lattice under componentwise order (D. Gale and L. Shapley were already aware of the existence of intermediate stable sets of marriages)—depending on conventions, the minimal element is the best possible stable matching for the men (obtained by running the Gale-Shapley algorithm with men proposing), while the maximal element is the best for the women. Although we are not aware of an equivalence between our work and the theory of stable marriages, we note that there is a certain similarity between the roughly symmetric Algorithms 3 and 5 and the Gale-Shapley proposing algorithms.[4] Our leftmost equitable partition and rightmost equitable partition are analogous to the male- and female-favoring stable sets of marriages.

6. Applications

6.1. **The Sweep Map.** By taking m sufficiently large, the modular sweep map emulates the sweep map of [ALW14b], as we now explain.

Fix $a := (a_1, \ldots, a_n) \in \mathbb{Z}^n$, let $e := (e_1, \ldots, e_n) \in \mathbb{N}^n$, and define $\mathcal{A}_\mathbb{Z}$ to be the set of words containing e_j copies of a_j for $1 \leq j \leq n$. For a word $\mathsf{w} = \mathsf{w}_1\mathsf{w}_2 \cdots \mathsf{w}_N \in \mathcal{A}_\mathbb{Z}$, define the *level* of w_j to be the integer $\ell_j := \sum_{i=1}^{j} \mathsf{w}_i$ for $1 \leq j \leq N$.

The *sweep map* is the function $\mathsf{sweep} : \mathcal{A}_\mathbb{Z} \to \mathcal{A}_\mathbb{Z}$ that sorts $\mathsf{w} \in \mathcal{A}_\mathbb{Z}$ according to its levels as follows: initialize $\mathsf{u} = \emptyset$ to be the empty word. For $k = -1, -2, -3, \ldots$ and then $k = \ldots, 3, 2, 1, 0$, read w from right to left and append to u all letters w_j whose level ℓ_j is equal to k. Define $\mathsf{sweep}(\mathsf{w}) := \mathsf{u}$.

Theorem 6.1 ([ALW14b, Conjecture 3.3 (a)])**.** *The sweep map is a bijection*

$$\mathsf{sweep} : \mathcal{A}_\mathbb{Z} \to \mathcal{A}_\mathbb{Z}.$$

Proof. Since the modular sweep map only *permutes* its input, it restricts to a bijection on words with a specified content. We claim that by choosing m large enough, the modular sweep map agrees with the sweep map when the letters a_j with multiplicities e_j are replaced by their natural representatives $a_j \pmod m$ in $\{0, 1, 2, \ldots, m-1\}$ (and all other elements are given multiplicity 0), and similarly for the levels ℓ_j.

When computing the modular sweep map, so long as m is large enough that all distinct letters correspond to distinct letters modulo m, and that all distinct levels correspond to distinct modular levels, the images of the modular sweep map and the sweep maps agree modulo m. We can achieve this simultaneously for all words in $\mathcal{A}_\mathbb{Z}$ by taking $m > \sum_{j=1}^{n} e_j|a_j|$. We conclude that by taking m large enough and the letters a_i modulo m, we may use the inverse modular sweep map to compute the inverse sweep map. □

[4] It is not too difficult to imagine a somewhat scandalous rephrasing of Section 1.1 using a polyamorous cul-de-sac.

Remark 6.2. It might seem that the definitions above are a weaker formulation than the one used in [ALW14b], which uses an arbitrary alphabet (not necessarily $\mathbb{Z}$) along with a weight function from the alphabet to $\mathbb{Z}$. In that language, our formulation has replaced the alphabet by its image under the weight function, so that it might seem that we have restricted ourselves only to the case of injective weight functions. In fact, the two formulations are equivalent for the following two reasons:

(1) the remembering map depends only on the weight;
(2) the inverse modular presweep map traverses the letters in a specified order.

To recover the formulation using arbitrary alphabets and weight functions is simple: use the remembering map on the image of the word by the weight function and then apply the inverse modular presweep map recording the letter from the alphabet (rather than its image under the weight function).

6.2. **Sweeping Dyck Words.** Preserving the notation of the previous section, let $\mathcal{A}_{\mathbb{N}} \subseteq \mathcal{A}_{\mathbb{Z}}$ be the subset of words in $\mathcal{A}_{\mathbb{Z}}$ whose levels are all nonnegative. Following [ALW14b], we call $\mathcal{A}_{\mathbb{N}}$ the set of *Dyck words*.

Theorem 6.3 ([ALW14b, Conjecture 3.3 (b)])**.** *The sweep map is a bijection*

$$\mathsf{sweep} : \mathcal{A}_{\mathbb{N}} \to \mathcal{A}_{\mathbb{N}}.$$

We shall interpret elements of $\mathcal{A}_{\mathbb{N}}$ as paths in the lattice $\mathbb{Z}^n$, with standard basis denoted $\mathsf{a}_1, \mathsf{a}_2, \ldots, \mathsf{a}_n$. By reading the letter a_i as a step in direction $\mathsf{a}_i = \mathsf{dir}(a_i)$, a word $\mathsf{w} \in \mathcal{A}_{\mathbb{N}}$ is interpreted as a lattice path $\mathsf{path}(\mathsf{w})$ from $(0, 0, \ldots, 0)$ to $(e_1, e_2, \ldots, e_n)$ using steps a_i, such that $x \cdot a \geq 0$ for any x on $\mathsf{path}(\mathsf{w})$. The *endpoint* of the step corresponding to w_j is the point $\sum_{i=1}^{j} \mathsf{dir}(\mathsf{w}_i)$. One can now visualize the sweep map as the process of *sweeping* the affine hyperplane defined by $\mathcal{H}_{a,k} := \{x : x \cdot a = k\}$ down from $k = \infty$ to $k = 0$ and recording the letters corresponding to the endpoints of the steps of $\mathsf{path}(\mathsf{w})$ in the order they are intersected by this hyperplane. Figure 2 illustrates this geometric interpretation for $n = 2$ with $a = (-7, 4)$ and $e = (4, 7)$.

This definition is slightly misleading—in order to have agreement with sweep, one must respect the tiebreaking "right to left" condition when two points are intersected by $\mathcal{H}_{a,k}$ *at the same time*. It is therefore nicer in this geometric interpretation to break these ties by slightly perturbing the vector a, so that we allow k to vary *continuously* from $k = \infty$ to $k = 0$ (and still record the letters corresponding to the endpoints of $\mathsf{path}(\mathsf{w})$ in the order they are intersected by $\mathcal{H}_{a,k}$). To break ties, we perturb $a = (a_1, \ldots, a_n)$ by choosing $a' = (a'_1, a'_2, \ldots, a'_n)$ linearly independent over $\mathbb{Q}$ such that $a'_i > a_i$ for all i, and such that there is no lattice point x in the hypercube spanned by $\{e_i \mathsf{a}_i\}_{i=1}^{n}$ such that $x \cdot a < 0$ and $x \cdot a' > 0$.

The following argument was suggested by M. Thiel, generalizing [ALW14b, Proposition 3.2]. (C. Reutenauer has developed a similar argument [Reu].)

Proof. Let $\mathsf{u} = \mathsf{sweep}(\mathsf{w}_1 \mathsf{w}_2 \cdots \mathsf{w}_N)$. We show that for any j, the initial segment $\mathsf{u}_1 \mathsf{u}_2 \cdots \mathsf{u}_j$ ends on or above $\mathcal{H}_{a,0}$, from which we may conclude that $\mathsf{u} \in \mathcal{A}_{\mathbb{N}}$. Fix $1 \leq j \leq N$. By definition of sweep, for any such j, there exists a k for which all letters of w arranged into this initial segment of u have levels greater than or equal to k. We visualize this as those steps of $\mathsf{path}(\mathsf{w})$ whose endpoint lies on or above $\mathcal{H}_{a',k}$. Translating to the origin, by construction, each of the connected pieces separately satisfy the Dyck word condition, so that their rearrangement ends on or above $\mathcal{H}_{a',0}$. Since we have chosen a' so that the sweep map adds letters one at a time, this holds for any j. □

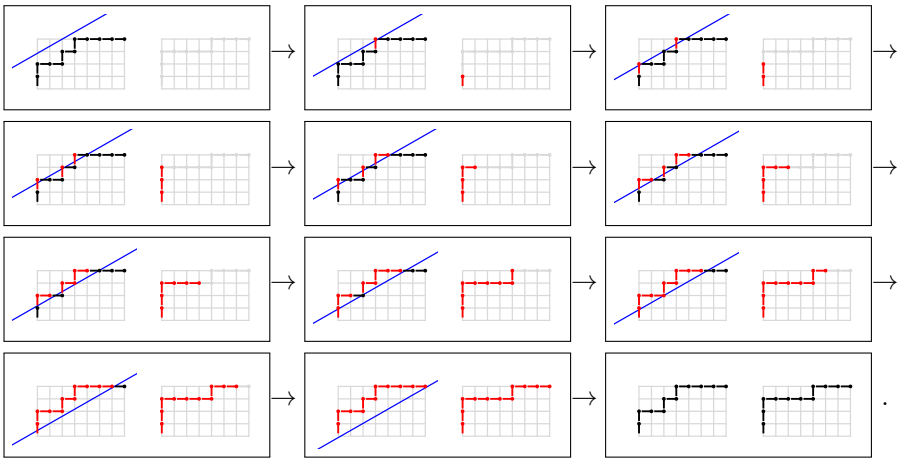

FIGURE 2. An illustration of the geometric interpretation of sweep. To form the right path, the steps of the left path are rearranged according to the order in which they are encountered by a line of slope 4/7 sweeping down from above.

6.3. **The Zeta Map.** Finally, we consider the special case of Dyck words for an alphabet $\{a, b\}$ of size $n = 2$, such that $a > 0$ and $b < 0$ and where the letter a occurs $-b$ times and the letter b occurs a times. We shall write this set of Dyck words as $\mathcal{D}_{a,b}$—these paths are of fundamental importance for the study of rational (type A) Catalan combinatorics [ARW13, ALW14a, GMV16, BR15, CDH16, Sul15, Xin15, GX16].

By [ALW14b, Table 1] and [ALW14b, Theorem 4.8, Lemma 4.10, Theorem 4.12], the *zeta map* may be defined as a variant of the sweep map $\zeta : \mathcal{D}_{a,b} \to \mathcal{D}_{a,b}$ that sorts $\mathsf{w} \in \mathcal{D}_{a,b}$ as follows: initialize $\mathsf{u} = \emptyset$ to be the empty word. For $k = 0, 1, 2, \ldots$ and then $k = \ldots, -3, -2, -1$, read w from left to right and append to u all letters w_j whose level ℓ_j is equal to k. Define $\zeta(\mathsf{w}) := \mathsf{u}$.

We have the following corollary of Theorem 6.1, which is of independent interest.

Corollary 6.4 (Zeta for Rational Dyck Paths)**.** *The zeta map is a bijection*

$$\zeta : \mathcal{D}_{a,b} \to \mathcal{D}_{a,b}.$$

Proof. For $\mathsf{w} = \mathsf{w}_1\mathsf{w}_2 \cdots \mathsf{w}_N$, let

$$\mathsf{rev}(\mathsf{w}) := \mathsf{w}_N \cdots \mathsf{w}_2\mathsf{w}_1 \quad \text{and}$$
$$-\mathsf{w} := (-\mathsf{w}_1)(-\mathsf{w}_2)\cdots(-\mathsf{w}_N).$$

Then the zeta map may be computed as

$$\zeta(\mathsf{w}) = -\left(\mathsf{rev} \circ \mathsf{sweep} \circ \mathsf{rev}\right)(-\mathsf{w}).$$

Since sweep is a bijection, we conclude that ζ is a bijection. □

Remark 6.5. When a and b are relatively prime, there cannot be any proper balanced block suffixes—that is, the only possible block suffix is the entire word. In this case, the distributive lattice of equitable partitions is just a chain, and every equitable partition is a cyclic shift of the blocks of the successful partition.

References

[AKOP02] George Andrews, Christian Krattenthaler, Luigi Orsina, and Paolo Papi, *ad-nilpotent $\mathfrak{b}$-ideals in $sl(n)$ having a fixed class of nilpotence: combinatorics and enumeration*, Transactions of the American Mathematical Society **354** (2002), no. 10, 3835–3853.

[ALW14a] Drew Armstrong, Nicholas Loehr, and Gregory Warrington, *Rational parking functions and Catalan numbers*, arXiv preprint arXiv:1403.1845 (2014).

[ALW14b] ———, *Sweep maps: A continuous family of sorting algorithms*, arXiv preprint arXiv:1406.1196 (2014).

[ARW13] Drew Armstrong, Brendon Rhoades, and Nathan Williams, *Rational associahedra and noncrossing partitions*, The Electronic Journal of Combinatorics **20** (2013), no. 3, P54.

[BR15] Michelle Bodnar and Brendon Rhoades, *Cyclic sieving and rational Catalan theory*, arXiv preprint arXiv:1510.08502 (2015).

[BWZ11] Chris Berg, Nathan Williams, and Mike Zabrocki, *Symmetries on the lattice of k-bounded partitions*, Annals of Combinatorics (2011), 1–31.

[CDH16] Cesar Ceballos, Tom Denton, and Christopher Hanusa, *Combinatorics of the zeta map on rational Dyck paths*, Journal of Combinatorial Theory, Series A **141** (2016), 33–77.

[Egg03] Eric Egge, *A Schröder generalization of Haglund's statistic on Catalan paths*, The Electronic Journal of Combinatorics **10** (2003), no. 1, 21.

[GH96] Adriano Garsia and Mark Haiman, *A remarkable q,t-Catalan sequence and q-Lagrange inversion*, Journal of Algebraic Combinatorics **5** (1996), no. 3, 191–244.

[GH02] Adriano Garsia and James Haglund, *A proof of the q,t-Catalan positivity conjecture*, Discrete Mathematics **256** (2002), no. 3, 677–717.

[GM14] Evgeny Gorsky and Mikhail Mazin, *Compactified Jacobians and q,t-Catalan numbers, II*, Journal of Algebraic Combinatorics **39** (2014), no. 1, 153–186.

[GMV16] Eugene Gorsky, Mikhail Mazin, and Monica Vazirani, *Affine permutations and rational slope parking functions*, Transactions of the American Mathematical Society (2016).

[GS62] David Gale and Lloyd S Shapley, *College admissions and the stability of marriage*, The American Mathematical Monthly **69** (1962), no. 1, 9–15.

[GX16] Adriano Garsia and Guoce Xin, *Inverting the rational sweep map*, arXiv preprint arXiv:1602.02346 (2016).

[Hag03] James Haglund, *Conjectured statistics for the q,t-Catalan numbers*, Advances in Mathematics **175** (2003), no. 2, 319–334.

[Knu76] Donald Knuth, *Mariages stables et leurs relations avec d'autres problèmes combinatoires: introduction à l'analyse mathématique des algorithmes*, Montréal: Presses de l'Université de Montréal, 1976.

[LLL14] Kyungyong Lee, Li Li, and Nicholas Loehr, *Combinatorics of certain higher q,t-Catalan polynomials: chains, joint symmetry, and the Garsia–Haiman formula*, Journal of Algebraic Combinatorics **39** (2014), no. 4, 749–781.

[Loe03] Nicholas Loehr, *Multivariate analogues of Catalan numbers, parking functions, and their extensions*, University of California, San Diego, 2003.

[Reu] Christophe Reutenauer, personal communication.

[Sul15] Robin Sulzgruber, *Rational Shi tableaux and the skew length statistic*, arXiv preprint arXiv:1512.04320 (2015).

[Sut02] Ruedi Suter, *Young's lattice and dihedral symmetries*, European Journal of Combinatorics **23** (2002), no. 2, 233–238.

[Sut04] ———, *Abelian ideals in a Borel subalgebra of a complex simple Lie algebra*, Inventiones mathematicae **156** (2004), no. 1, 175–221.

[TW14] Hugh Thomas and Nathan Williams, *Cyclic symmetry of the scaled simplex*, Journal of Algebraic Combinatorics **39** (2014), no. 2, 225–246.

[VW12] Mirko Visontai and Nathan Williams, *Dendrodistinctivity*, GASCom 2012 (2012).

[Wil96a] Kevin Jon Williams, *Examining the NRMP algorithm*, Academic Medicine **71** (April, 1996), no. 4, 310–312.

[Wil96b] ———, *Match algorithms revisited*, Academic Medicine **71** (May, 1996), no. 5, 415–416.

[Wil11] Nathan Williams, *Bijactions*, Master's thesis, University of Minnesota (2011).

[WWW81] Kevin Jon Williams, Victoria Werth, and Jon Wolff, *An analysis of the resident match*, New England Journal of Medicine **304** (1981), no. 19, 1165–1166, Correspondence in ——— **305** (1981), no.9, p.526.

[Xin15] Guoce Xin, *An efficient search algorithm for inverting the sweep map on rational Dyck paths*, arXiv preprint arXiv:1505.00823 (2015).

Appendix A. Differences with [TW14]

Although Algorithms 1 and 2 in [TW14] are *very slightly* different from Algorithms 1 and 2 in this paper, the notion of a successful partition here and in [TW14] agree, as we now explain.

There are three differences between Algorithm 2 here and Algorithm 2 in [TW14]. The first difference is the letters recorded—at the ith step, rather than recording the first letter of the current block ℓ_{N-i+1} (as we do here), in [TW14] we instead record the label of the next block ℓ_{N-i}.

The second and third differences are purely notational. We labeled blocks in [TW14, Definition 7.3] as

$$\mathsf{u}^*_{\ell_N+1}|\mathsf{u}^*_{\ell_N+2}|\cdots|\mathsf{u}^*_{\ell_N},$$

while here (for consistency with [ALW14b]) we label them as

$$\mathsf{u}^*_{m-1}|\mathsf{u}^*_{m-2}|\cdots|\mathsf{u}^*_0.$$

Algorithm 2 here compensates for this by beginning with block $\mathsf{u}^*_{\ell_N}$ (as opposed to u^*_0) and subtracting (as opposed to adding) the first letter of $\mathsf{u}^*_{\ell_{N-i+1}}$ when computing ℓ_{N-i}.

Finally, what we call the "rightmost equitable partition" in [TW14] has now become the "leftmost equitable partition." In [TW14], we chose the terminology "rightmost" to represent the positions of the *block dividers* in the word u, rather than the positions of the letters of u in the blocks—we have changed this convention here to be more intuitive.

(H. Thomas) LaCIM, Université de Québec à Montéal, Montréal (Québec), Canada
E-mail address: hugh.ross.thomas@gmail.com

(N. Williams) LaCIM, Université de Québec à Montéal, Montréal (Québec), Canada
E-mail address: nathan.f.williams@gmail.com